

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	471	((symmetric adj2 multiprocessor\$5) smp) with (shar\$5 adj2 memor\$5)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 17:01
L2	282	1 and @ad<"20010504"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 17:01
L3	50	((symmetric adj2 multiprocessor\$5) smp) with interconnect\$5 with (shar\$5 adj2 memor\$5)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 17:14
L4	32	3 and @ad<"20010504"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 17:15
L5	5	((symmetric adj2 multiprocessor\$5) smp) with bandwidth\$5 with (shar\$5 adj2 memor\$5)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 17:04
L6	72	((symmetric adj2 multiprocessor\$5) smp) and (bandwidth\$5 with (shar\$5 adj2 memor\$5))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 17:05
L7	47	6 and @ad<"20010504"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 17:07
L11	1736	(high adj2 bandwidth adj2 (link\$3 channel\$3 line\$3 interconnect\$5))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 17:13

L12	6621	((symmetric adj2 multiprocessor\$5) smp)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 17:14
L13	20070	shar\$5 adj memor\$5	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 17:14
L14	26	11 and 12	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 17:14
L15	42214	14 amd 13	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 17:14
L16	13	11 and 12 and 13	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 17:15
L17	6	16 and @ad<"20010504"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 17:25
L22	2	("6370145").PN.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 17:55
S1	2	("6587875").PN.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/23 20:16

S2	1262	(link\$3 medium) near9 share\$3 near9 memor\$5	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/23 20:17
S3	500	(link\$3 medium) near9 share\$3 adj memor\$5	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 09:59
S4	301	S3 and @ad<"20010504"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 09:59
S5	8	(S3 with bandwidth\$5) and @ad<"20010504"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 09:59
S6	10	(S3 same bandwidth\$5) and @ad<"20010504"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/23 20:28
S7	86	(S3 and bandwidth\$5) and @ad<"20010504"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 09:37
S11	75	share\$3 near memor\$5 near link\$3	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 09:50
S12	46	S11 and @ad<"20010504"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 09:50

S13	934	(link\$3 channel\$3 medium) near9 share\$3 adj memor\$5	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 10:39
S14	537	S13 and @ad<"20010504"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 17:01
S17	1167	(link\$3 channel\$3 interconnect\$6 medium) near9 share\$3 adj memor\$5	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 10:16
S19	261	(link\$3 channel\$3 interconnect\$5 medium) near share\$3 near memor\$5	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 10:41
S20	161	S19 and @ad<"20010504"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 10:41
S21	113	share\$3 adj memor\$5 adj (link\$3 channel\$3 interconnect\$5 medium)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 10:41
S22	62	S21 and @ad<"20010504"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 10:41
S25	351	(interconnect\$5 adj3 shar\$6 adj2 (medium memor\$5 channel\$5 connection\$5 link\$3 line\$3))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 12:56

S26	347	(interconnect\$5 adj3 shar\$6 adj2 (medium memor\$5 channel\$5 connection? link\$3 line\$3))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 12:56
S27	261	S26 and @ad<"20010504"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 12:58
S28	3	(S26 with bandwidth\$5) and @ad<"20010504"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 12:59
S29	13	(S26 same bandwidth\$5) and @ad<"20010504"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 13:03
S30	2665	(link\$3 channel\$5 line\$3 interconnect\$5 connect connects connection connecting) near7 shar\$5 adj (memor\$5 medium\$5)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 13:03
S31	63	(S30 same bandwidth\$5) and @ad<"20010504"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 13:32
S32	1675	S30 and @ad<"20010504"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 14:16
S33	81	(S30 with multiprocessor\$5) and @ad<"20010504"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 13:33

S35	25	multiprocessor\$5 with (shar\$3 adj (connection\$3 interconnect\$5 connecting\$3 link\$3 line\$3 channel\$3))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 14:16
S36	11	S35 and @ad<"20010504"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/06/24 17:00


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☐ The ACM Digital Library ☒ The Guide

THE GUIDE TO COMPUTING LITERATURE

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

The Mercury Interconnect Architecture: a cost-effective infrastructure for high-performance servers

Full text Pdf (1.53 MB)

 Source **ACM SIGARCH Computer Architecture News** [archive](#)

 Volume 25, Issue 2 (May 1997) [table of contents](#)

Pages: 98 - 107

Year of Publication: 1997

ISBN:0-89791-901-7

Also published in ...

Authors [Wolf-Dietrich Weber](#) HAL Computer Systems, 1315 Dell Ave, Campbell, CA
[Stephen Gold](#) HAL Computer Systems, 1315 Dell Ave, Campbell, CA
[Pat Helland](#) Microsoft Corporation, One Microsoft Way, Redmond, WA
[Takeshi Shimizu](#) HAL Computer Systems, 1315 Dell Ave, Campbell, CA
[Thomas Wicki](#) HAL Computer Systems, 1315 Dell Ave, Campbell, CA
[Winfried Wilcke](#) HAL Computer Systems, 1315 Dell Ave, Campbell, CA

Publisher ACM Press New York, NY, USA

 Additional Information: [abstract](#) [references](#) [citations](#) [index terms](#) [collaborative colleagues](#) [peer to peer](#)

Tools and Actions: [Discussions](#) [Find similar Articles](#) [Review this Article](#)
[Save this Article to a Binder](#) Display Formats: [BibTex](#) [EndNote](#) [ACM Ref](#)

DOI Bookmark: Use this link to bookmark this Article: <http://doi.acm.org/10.1145/384286.264149>
[What is a DOI?](#)

↑ ABSTRACT

This paper presents HAL's Mercury Interconnect Architecture, an interconnect infrastructure designed to link commodity microprocessors, memory, and I/O components into high-performance multiprocessing servers. Both shared-memory and message-passing systems, as well as hybrid systems are supported by the interconnect. The key attributes of the Mercury Interconnect Architecture are: low latency, high bandwidth, a modular and flexible design, reliability/availability/serviceability (RAS) features, and a simplicity that enables very cost-effective implementations. The first implementation of the architecture links multiple 4-processor Pentium[®] Pro based nodes. In a 4-node (16-processor) shared-memory configuration, this system achieves a remote read latency of just over 1 μ s, and a maximum interconnect bandwidth of 6.4 GByte/s. Both of these parameters far outpace comparable SCI-based solutions, while utilizing much fewer hardware components.

↑ REFERENCES

Note: OCR errors may be found in this Reference List extracted from the full text article. ACM has opted to expose the complete List rather than only correct and linked references.

1 L. Censier and P. Feautrier. A New Solution to Coherence Problems in Multicache Systems. IEEE Transactions on Computers C-27, 12 (December 1978), pages 1112-1118.

2 R. Clark and K. Alnes. An SCI Interconnect Chipset and Adapter. In Symposium Record, Hot

Interconnects 1V, pages 221-235, August 1996.

3 IEEE Computer Society. IEEE Standard for Scalable Coherent Interface (SCI), IEEE Standard 1596-1992, New York, August 1993.

4 Intel. Pentium~ Pro Family Developer's Manual Volume 3: Operating System Writer's Guide. Intel Corporation, 1996.

5 D. Lenoski, J. Laudon, K. Gharaehorloo, A. Gupta, and J. Hennessy. The Directory-Based Cache Coherence Protocol for the DASH Multiprocessor. In Proceedings of the 17th International Symposium on Computer Architecture, pages " 148-159, May 1990.

6 Daniel Lenoski , James Laudon , Kourosh Gharachorloo , Wolf-Dietrich Weber , Anoop Gupta , John Hennessy , Mark Horowitz , Monica S. Lam, The Stanford Dash Multiprocessor, Computer, v.25 n.3, p.63-79, March 1992

7 Daniel E. Lenoski , Wolf-Dietrich Weber, Scalable Shared-Memory Multiprocessing, Morgan Kaufmann Publishers Inc., San Francisco, CA, 1994

8 Tom Lovett , Russell Clapp, STING: a CC-NUMA computer system for the commercial marketplace, Proceedings of the 23rd annual international symposium on Computer architecture, p.308-317, May 22-24, 1996, Philadelphia, Pennsylvania, United States

9 T. Lovett and R. Clapp. Private correspondence, February 1997.

10 A. Mu, B. Chia, S. Kondapalli, C. Koo, J. Larson, L. Nguyen, R. Sastry, Y. Satsukawa, H.-C. Shih, T. Wicki, C. Wu, K. Yu, and X. Zhang. A 285 MHz 6-port Plesiochronous Router Chip with Non-Blocking Cross-Bar Switch. In 1996 Symposium on VLSI Circuits: Digest of Technical Papers, pages 136-137, 1996.

11 Albert Mu , Jeff Larson , Raghu Sastry , Thomas Wicki , Winfried W. Wilcke, A 9.6 GigaByte/s Throughput Plesiochronous Routing Chip, Proceedings of the 41st IEEE International Computer Conference, p.261, February 25-28, 1996

12 W.-D. Weber. Scalable Directories for Cache-Coherent Shared-Memory Multiprocessors. Ph.D. Thesis, Stanford University, January 1993. Also available as Stanford University Technical Report CSL-TR-93-557.

↑ CITINGS 10

Donglai Dai , Dhabaleswar K. Panda, Exploiting the Benefits of Multiple-Path Network in DSM Systems: Architectural Alternatives and Performance Evaluation, IEEE Transactions on Computers, v.48 n.2, p.236-244, February 1999

Chris Wilson , David L. Dill, Reliable verification using symbolic simulation with scalar values, Proceedings of the 37th conference on Design automation, p.124-129, June 05-09, 2000, Los Angeles, California, United States

An-Chow Lai , Babak Falsafi, Selective, accurate, and timely self-invalidation using last-touch prediction, ACM SIGARCH Computer Architecture News, v.28 n.2, p.139-148, May 2000

An-Chow Lai , Babak Falsafi, Comparing the effectiveness of fine-grain memory caching against page migration/replication in reducing traffic in DSM clusters, Proceedings of the twelfth annual ACM symposium on Parallel algorithms and architectures, p.79-88, July 09-13, 2000, Bar Harbor, Maine, United States

Jaeheon Jeong , Michel Dubois, Optimal replacements in caches with two miss costs, Proceedings of

the eleventh annual ACM symposium on Parallel algorithms and architectures, p.155-164, June 27-30, 1999, Saint Malo, France

An-Chow Lai , Babak Falsafi, Memory sharing predictor: the key to a speculative coherent DSM, ACM SIGARCH Computer Architecture News, v.27 n.2, p.172-183, May 1999

Dan Teodosiu , Joel Baxter , Kinshuk Govil , John Chapin , Mendel Rosenblum , Mark Horowitz, Hardware fault containment in scalable shared-memory multiprocessors, ACM SIGARCH Computer Architecture News, v.25 n.2, p.73-84, May 1997

Mark Heinrich , Vijayaraghavan Soundararajan , John Hennessy , Anoop Gupta, A Quantitative Analysis of the Performance and Scalability of Distributed Shared Memory Cache Coherence Protocols, IEEE Transactions on Computers, v.48 n.2, p.205-217, February 1999

Tao Li , Lizy Kurian John, ADir pNB: A Cost-Effective Way to Implement Full Map Directory-Based Cache Coherence Protocols, IEEE Transactions on Computers, v.50 n.9, p.921-934, September 2001

↑ INDEX TERMS

Primary Classification:

C. Computer Systems Organization

↳ C.1 PROCESSOR ARCHITECTURES

↳ C.1.2 Multiple Data Stream Architectures (Multiprocessors)

↳ **Subjects:** Interconnection architectures (e.g., common bus, multiport memory, crossbar switch)

Additional Classification:

C. Computer Systems Organization

F. Theory of Computation

↳ F.2 ANALYSIS OF ALGORITHMS AND PROBLEM COMPLEXITY

↳ F.2.2 Nonnumerical Algorithms and Problems

↳ **Subjects:** Routing and layout

General Terms:

Design, Measurement, Performance, Reliability, Theory

↑ Collaborative Colleagues:

Stephen Gold: Pat Helland
 Takeshi Shimizu
 Wolf-Dietrich Weber
 Thomas Wicki
 Winfried Wilcke

<u>Pat Helland:</u>	<u>Larry Cable</u>	<u>Andreas Reuter</u>
	<u>Richard Carr</u>	<u>Harald Sammer</u>
	<u>Matthieu Devin</u>	<u>Dennis Shasha</u>
	<u>Scott Dietzen</u>	<u>Takeshi Shimizu</u>
	<u>Phil Garrett</u>	<u>Wolf-Dietrich Weber</u>
	<u>Stephen Gold</u>	<u>Thomas Wicki</u>
	<u>Jim Gray</u>	<u>Winfried Wilcke</u>
	<u>Jim Lyon</u>	<u>Dan Wolfson</u>
	<u>C. Mohan</u>	
	<u>Patrick O'Neil</u>	

Takeshi Shimizu: John Boreczky

[Stephen Gold](#)
[Pat Helland](#)
[Stephen W. Smoliar](#)
[Wolf-Dietrich Weber](#)
[Thomas Wicki](#)
[Winfried Wilcke](#)

[Wolf-Dietrich Weber](#): [Kourosh Gharachorloo](#) [Todd Mowry](#)
[Stephen Gold](#) [Takeshi Shimizu](#)
[Anoop Gupta](#) [Jaswinder Pal Singh](#)
[Pat Helland](#) [Thomas Wicki](#)
[John Hennessy](#) [Winfried Wilcke](#)
[Mark Horowitz](#)
[Monica S. Lam](#)
[James Laudon](#)
[Daniel Lenoski](#)
[Daniel E. Lenoski](#)

[Thomas Wicki](#): [Stephen Gold](#)
[Pat Helland](#)
[Jeff Larson](#)
[Albert Mu](#)
[Raghu Sastry](#)
[Takeshi Shimizu](#)
[Wolf-Dietrich Weber](#)
[Winfried Wilcke](#)
[Winfried W. Wilcke](#)

[Winfried Wilcke](#): [Robert Garner](#)
[Stephen Gold](#)
[Pat Helland](#)
[Takeshi Shimizu](#)
[Wolf-Dietrich Weber](#)
[Thomas Wicki](#)

↑ Peer to Peer - Readers of this Article have also read:





- [Data structures for quadtree approximation and compression](#)
Communications of the ACM 28, 9
Hanan Samet
- [A hierarchical single-key-lock access control using the Chinese remainder theorem](#)
Proceedings of the 1992 ACM/SIGAPP Symposium on Applied computing
Kim S. Lee , Huizhu Lu , D. D. Fisher
- [The GemStone object database management system](#)
Communications of the ACM 34, 10
Paul Butterworth , Allen Otis , Jacob Stein
- [Putting innovation to work: adoption strategies for multimedia communication systems](#)
Communications of the ACM 34, 12
Ellen Francik , Susan Ehrlich Rudman , Donna Cooper , Stephen Levine
- [An intelligent component database for behavioral synthesis](#)
Proceedings of the 27th ACM/IEEE conference on Design automation
Gwo-Dong Chen , Daniel D. Gajski

↑ This Article has also been published in:

- [International Conference on Computer Architecture](#)
Proceedings of the 24th annual international symposium on Computer architecture
1997 , Denver, Colorado, United States

The ACM Portal is published by the Association for Computing Machinery. Copyright ?2005 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)